Perpetual Codes for Network Coding

Janus Heide*, Morten V. Pedersen*, Frank H.P. Fitzek* and Muriel Médard†
* Aalborg University, Aalborg, Denmark, Email: [jah|mvp|ff]@es.aau.dk
† Massachusetts Institute of Technology, Cambridge, Massachusetts, Email: medard@mit.edu

Abstract—Random Linear Network Coding (RLNC) provides a theoretically efficient method for coding. Some of its practical drawbacks are the complexity of decoding and the overhead due to the coding vectors. For computationally weak and battery-driven platforms, these challenges are particular important. In this work, we consider the coding variant *Perpetual codes* which are sparse, non-uniform and the coding vectors have a compact representation. The sparsity allows for fast encoding and decoding, and the non-uniform protection of symbols enables recoding where the produced symbols are indistinguishable from those encoded at the source. The presented results show that the approach can provide a coding overhead arbitrarily close to that of RLNC, but at reduced computational load. The achieved gain over RLNC grows with the generation size, and both encoding and decoding throughput is approximately one order of magnitude higher compared to RLNC at a generation size of 2048. Additionally, the approach allows for easy adjustment between coding throughput and code overhead, which makes it suitable for a broad range of platforms and applications.

Index Terms—Network coding, implementation, algorithms, complexity

arXiv:1509.04492v1 [cs.NI] 15 Sep 2015

1 Introduction

Network Coding (NC) is a promising paradigm [1] that has been shown to provide benefits in many different networks and applications. NC enables coding at individual nodes in a communication network, and thus is fundamentally different from the end-to-end approach of channel and source coding. With NC, packets are no longer treated as atomic entities since they can be combined and re-combined at any node in the network. This allows for a less restricted view on the flow of information in networks, which can be particularly helpful when building distribution systems for less structured networks such as meshed, peer-to-peer or highly mobile networks.

In this work, we focus on random NC approaches, i.e. RLNC [2], and disregard deterministic coding. The reason is that our primary interest is cooperative and highly mobile wireless networks, which fit perfectly with the highly decentralized nature of RLNC. In particular, RLNC reduces the signaling overhead and increases robustness towards changing channel conditions in the network. At the same time, it allows for the construction of much simpler distribution systems, which is desirable from an engineering point of view.

Unfortunately, RLNC is inherently computationally demanding that has spawned several efforts to produce optimized implementations and modify the underlying code [3], [4]. Even though several solutions and implementations have been declared to provide *sufficient coding throughput* continued efforts are valid as they can ensure higher coding throughput. Computational resources can be conserved tasks, such as video decoding, and the energy consumption introduced by coding can be reduced further. This is of particular importance when NC is deployed on battery-driven devices with modest computational capabilities.

This paper presents our work on applying Perpetual Codes, which was suggested and named in the unpublished draft [5], for NC systems. The encoding is sparse and non-uniform which allows for fast decoding as *fill-in* [6] is avoided while recoding is still possible. The approach presented here is similar to what is called a *smooth perpetual code* in [5], but with two significant differences, neither zero padding nor a pre-code is used. This simplifies the analysis, but complicates the final decoding step. We describe how encoding and decoding can be performed and analyze the overhead and the complexity. We verify our results with our own C++ implementation, which also provides practical throughput results. Furthermore, we describe, implement, and evaluate recoding which was not considered in [5]. The main insight from our results is that RLNC is a better choice at low to medium generation sizes, but perpetual codes are more suitable at medium to high sized generations. Hence, perpetual codes are not a substitution, but a supplement to RLNC.

This paper is primarily intended for researchers and developers who work with reliable data distribution on wireless and mobile platforms. Therefore we provide a short overview of RLNC and related work in Section 2. The approach to encoding, decoding and recoding is presented in Section 3 together with algorithms aimed at implementations in C or C++. Section 4 provides analysis of the performance of the code in terms of decoding complexity and code overhead and compares measurements results obtained from our implementation with the analytical expressions. Readers primarily interested in theoretical results and familiar with Forward Error Correction (FEC) and RLNC could therefore skip Section 2 as well as some parts of Section 3.

2 RLNC and Related Work

When data is distributed from one or more sources to one or more sinks using RLNC, then it is *encoded* at the sources to produce coded symbols and coding vectors that describe the encoding procedure. Together a coded symbol and a coding vector form a coded packet. When a sink has received *enough* coded packets, it can *decode* the original data. Additionally, received symbols can be recombined and thus *recoded*, at any relaying nodes in between the sources and sinks.

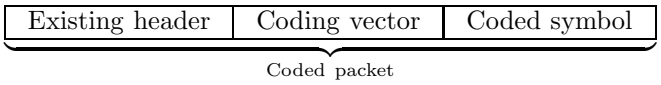

For practical reasons, the original data is typically divided into *generations* [7] of size g. We denote the data in such a generation as $\boldsymbol{M}$. This ensures that coding can be performed over data of any size, and that the performance of RLNC is independent of the data size. Each generation is divided into symbols, denoted as $\boldsymbol{m}_i$, and these symbols are then combined at random to create a coded symbol, denoted as $\boldsymbol{x}_i$. As all operations are performed over a Finite Field (FF) $\mathbb{F}_q$, the code is linear, and thus new valid coded symbols can be created from coded and non-coded symbols. Fig. 1 illustrates how the original symbols can be combined at random and provide an endless stream of coded symbols. The original data can be decoded by inverting the coding operations performed on the coded symbols. See [8], [9] for introductions to FF and RLNC.

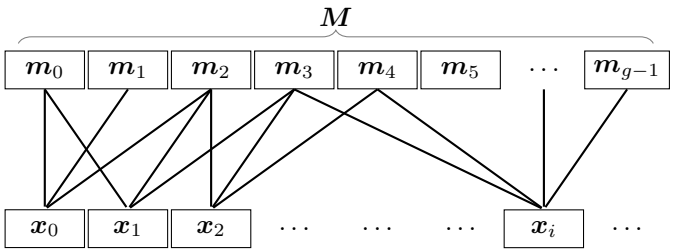

Fig. 1: Coded symbols are created from the original data.

Dividing the data into generations reduces both the computational work and the decoding delay. Unfortunately, it also introduces the need for additional signaling [10], [11], as each of the generations must be decoded successfully before the original data is recovered fully. It also increases the probability that the sink receives linearly dependent symbols which adds to the overhead of the code. This overhead is well understood for network typologies where (it can be assumed that) symbols are only received from sources that hold the original information [12]–[14]. In such systems, the parameters of the code can be chosen so that the overhead tends to zero and can be ignored. However, these parameters present a trade-off where higher values will generally result in lower code overhead but lower coding throughput [4]. The coding throughput also depends on less deterministic parameters, e.g. the hardware platform, programming language, and implementation optimizations [14]–[18]. Therefore a universally optimal set of values cannot be identified, as they depend on the system and on the target platform.

Some simplifications that can increase the coding throughput of RLNC are binary, systematic, and sparse variants [14], [19]–[21]. Binary codes are in widespread use and can obtain a low code overhead. They can be fast as operations in the binary field can be performed in parallel by all modern computers. Using a systematic code comes with no cost in term of overhead, and can potentially provide a high gain in both encoding and decoding throughput. Unfortunately, it is not possible to us this approach at every node, but only at the sources. Thus there is no or little gain if recoding is performed, which is the main reason to use RLNC in the first place. Using a sparse random code provides similar benefits and drawbacks as a systematic code. It becomes impractical to perform recoding, and the gain in decoding throughput can be small or non-existing [22].

Alternatively, the underlying code can be fundamentally modified or replaced to ensure a lower decoding complexity. A noteworthy suggestion is to use a convolutional code as the underlying code [23], [24] as they have been used in communication systems for many years. These efforts are still primarily theoretical as to the best of our knowledge currently no implementation of convolutional codes for NC exists. The work on perpetual codes [5] is related to this work, since it uses a related fundamental concept, combined with a concatenated approach similarly to Raptor codes [25]. However the authors aim was to propose a cache-friendly rateless erasure code, and they did not consider recoding, which is a necessary feature when used in a system that exploits NC. We note that linear block codes and convolutional codes may in some cases be *equivalent*, as they can describe codes with similar realizations using different terminology [26], [27].

Another direction in the search for improved trade-off between computational work and code overhead was suggested in [28]. Here the authors considered coding over several generations, called a *random annex* code [11], [29]. Each generation is extended to include symbols from other generations and thus when a generation is decoded these extra symbols are released. This reduces the problem of ensuring that all generations are decoded, and thus the overhead. At the same time, it is less computationally demanding as the decoding is performed in an inner and an outer step. The approach is very useful for file transfers, but less so for streaming as the final decoding delay is high as generations are not decoded sequentially. Additionally, the problem of how recoding could be performed has so far not been considered. Note that the idea of a *random annex* can be applied to many underlying codes, including the perpetual code considered in this work.

3 Code Operation

This section introduces the code and the three operations, *encoding*, *decoding* and *recoding*, that can be performed at nodes in the network. The notation used for analysis and algorithms is listed in Table 1. In vectors and matrices, we denote the first element with index zero. In some algorithms, the value -1 is used to denote non-valid or non-existing.

TABLE 1: Notation used for analysis and algorithms

Symbol	Definition
g	Generation size
q	Field size
w	Coding vector width
$\mathbb{F}_q$	A finite field with q elements
$\boldsymbol{g}$	Coding vector with g elements, starting at element 0
$\boldsymbol{G}$	Matrix containing all received coding vectors
$\boldsymbol{x}$	Coded symbol
$\boldsymbol{X}$	Matrix containing all received symbol
$\boldsymbol{h}$	Local recoding vector
$\boldsymbol{G}_i$	The ith row of the coding matrix
$\boldsymbol{G}_{i,j}$	The index in row i and column j of the coding matrix.
$\boldsymbol{X}_i$	The ith row of the symbol matrix
$\boldsymbol{X}_{i,j}$	The index in row i and column j of the symbol matrix.
p	Local variables for pivot indices
?	A randomly drawn integer.

In RLNC, the elements in the coding vector $\boldsymbol{g}$ are drawn completely at random, and thus each coded symbol is a combination of all the original symbols in one generation. This is not the case for the perpetual approach that we consider in this work. Instead, an element with index p is chosen as the pivot and the following w elements are drawn at random from $\mathbb{F}_q$. We denote w as the width of the coding vector. See Fig. 2 for a small example of some resulting coding vectors.

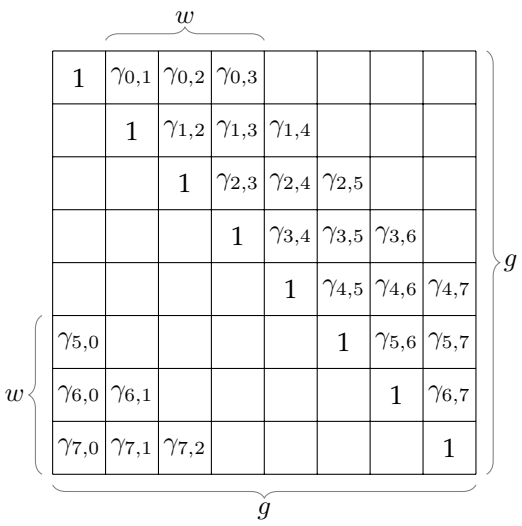

Fig. 2: All possible coding vectors, when $g = 8$ and $w = 3$. The γ's denote randomly drawn elements from $\mathbb{F}_q$.

3.1 Encoding

The data to be transmitted from the source is divided into generations, we denote the data in such a generation $\boldsymbol{M}$. Each generation is divided into g symbols that are represented with one or more FF elements in $\mathbb{F}_q$. The symbols are combined as specified by the coding vector $\boldsymbol{g}$ in order to create coded symbols $\boldsymbol{x}$.

$$\boldsymbol{x} = \boldsymbol{M} \cdot \boldsymbol{g} \tag{1}$$

The construction of a coding vector $\boldsymbol{g}$ and the corresponding coded symbol $\boldsymbol{x}$ is described by Algorithm 1.

Algorithm 1: encode

Input: $\boldsymbol{M}$
$\boldsymbol{g} \leftarrow \boldsymbol{0}$
$p \leftarrow (? \mod g)$
$\boldsymbol{g}_p \leftarrow 1$
for $i \in (p, p+w]$ **do**
  $\boldsymbol{g}_{(i \mod g)} \leftarrow (? \mod q)$
$\boldsymbol{x} \leftarrow \boldsymbol{M} \cdot \boldsymbol{g}$
return $\boldsymbol{g}, \boldsymbol{x}$

An index in the generation is drawn at random and used as the pivot, $p \in [0, g)$. The index in $\boldsymbol{g}$ that corresponds to this pivot element is set to one. For the subsequent w indices in $\boldsymbol{g}$, an element is drawn at random from $\mathbb{F}_q$. The remaining elements in $\boldsymbol{g}$ are zeros. The resulting coding vector is of the form illustrated in Fig. 2. To create a coded symbol, the coding vector is multiplied with the data, $\boldsymbol{x} = \boldsymbol{M} \cdot \boldsymbol{g}$. Together, the coding vector $\boldsymbol{g}$ and coded symbol $\boldsymbol{x}$ form a coded packet.

It is trivial to represent the coding vector in a very compact way. Each coding vector can be represented by an index and w scalars. The necessary bits for their

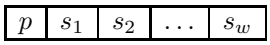

representation is given by Equation (2). The index can take g values and each of the w elements can take q values.

$$|\boldsymbol{g}| = \log_2(g) + w \cdot \log_2(q) \text{ [bits]} \tag{2}$$

Coding vectors can be generated in slightly different ways depending on how p is drawn and the size of w, see Table 2. The *systematic* mode does not produce coding vectors of the specified form, but we include it for completeness.

TABLE 2: Different encoding modes.

Mode	p drawn	w
Random	random $\in [0, g)$	$0 < w < g$
Sequential	sequentially looping from 0 to g-1	$0 < w < g$
Systematic	sequentially from 0 to g-1, subsequently drawn at random $\in [0, g)$	$w = 0$

3.2 Decoding

A node that receives coded packets can decode the original data by collecting the coded symbols in $\hat{\boldsymbol{X}}$ and the coding vectors in $\hat{\boldsymbol{G}}$. The original information, $\boldsymbol{M}$ can be found as in Equation (3), provided that $\hat{\boldsymbol{G}}$ is invertible and thus has full rank.

$$\boldsymbol{M} = \hat{\boldsymbol{X}} \cdot \hat{\boldsymbol{G}}^{-1} \tag{3}$$

To decode the original data in $\hat{\boldsymbol{M}}$, $\hat{\boldsymbol{G}}$ must be reduced to identity form by performing basic row operations that are simultaneously performed on $\hat{\boldsymbol{M}}$. When it is not possible to fully decode a symbol upon reception, then it is partially decoded, and stored for later processing. This is referred to as *on-the-fly* decoding. When enough symbols have been received so that $\hat{\boldsymbol{G}}$ has full rank, all received symbols can be fully decoded and the original data can be retrieved, we refer to this as *final* decoding.

Unlike RLNC and Sparse Random Linear Network Coding (SRLNC), this perpetual approach defines the location of the non-zero values in the coding vector. This makes it possible to decode symbols efficiently and without the problematic fill-in that can be observed during decoding and recoding SRLNC [22].

3.2.1 On-the-fly Decoding

When a new coded packet arrives, its coding vector is inserted into the decoding matrix iff. it has a *pivot candidate* that was not previously identified. We distinguish between *pivot* and *pivot candidate* as the element that is used as the pivot may only be found during the final decoding. Otherwise the previously received symbol with the same pivot candidate is subtracted from the new symbol, and the pivot candidate of the new symbol is changed. This is repeated until a new pivot candidate is identified. If the coding vector is reduced to the zero vector, the symbol is discarded.

In Fig. 3, we assume that three coded packets have been received and their coding vectors have been inserted into the decoding matrix. The pivot candidates of the received packets are zero, one, and seven, respectively. Subsequently, a coded packet with pivot candidate zero is received. This is denoted with a filled circle and arrow pointing to the coding vector of the packet in the left hand-side matrix. A symbol with the same pivot candidate have already been identified. Therefore, the existing row zero is subtracted from the incoming packet. This is denoted with the arrow pointing left into the left hand-side matrix. The element that initially was the pivot candidate is now zero and an element to the right has now become the pivot candidate. This step is repeated for the new pivot candidate and row one is subtracted from the incoming packet and element two becomes the pivot candidate. As this pivot candidate was previously not identified, the coding vector is inserted into the decoding matrix, which is marked with orange and the arrow pointing right into the decoding matrix.

A special case is when the on-the-fly phase causes the pivot candidate to wrap around to the start of the coding vector. An example of this is illustrated in Fig. 4. The incoming packet has pivot candidate seven for which a pivot candidate has already been identified in $\hat{\boldsymbol{G}}$. Thus row seven in $\hat{\boldsymbol{G}}$ is subtracted from the incoming packet. If the last element in the coding vector is reduced to the zero vector, the first element in the vector is considered next and becomes the pivot candidate. In this case, the resulting coding vector has a zero at index seven and thus the pivot candidate is now index zero. The packet is then further reduced similarly to the example in Fig. 3.

Algorithm 2: forwardSubstitute

Input: $\boldsymbol{g}, \boldsymbol{x}$
while $\boldsymbol{g} \neq \boldsymbol{0}$ **do**
  $p \leftarrow \text{pivot}(\boldsymbol{g})$
  if $\boldsymbol{G}_p \neq \boldsymbol{0}$ **then**
    $\boldsymbol{g} \leftarrow \boldsymbol{g} \cdot \frac{1}{\boldsymbol{g}_p} \oplus \boldsymbol{G}_p$
    $\boldsymbol{x} \leftarrow \boldsymbol{x} \cdot \frac{1}{\boldsymbol{g}_p} \oplus \boldsymbol{X}_p$ ▷ substitute into new packet
  else
    $\boldsymbol{G}_p \leftarrow \boldsymbol{g} \cdot \frac{1}{\boldsymbol{g}_p}$
    $\boldsymbol{X}_p \leftarrow \boldsymbol{x} \cdot \frac{1}{\boldsymbol{g}_p}$ ▷ insert new packet
    return p
return -1

In Algorithm 2, the existing row with the same pivot candidate is substituted into the received symbol, unless the received coding vector has been reduced to the zero vector. If a new pivot candidate is identified, then the coding vector and symbol are inserted into the respective matrices. Importantly, this algorithm guarantees that w will not increase during decoding.

The coding vector can be reduced to the zero vector if it is a linear combination of previously received coding vectors. It is possible to end in a dead-lock where a sequence of rows is repeatedly subtracted from the new packet. To avoid this, decoding should be terminated after some attempts and the packet discarded. From practical experiments, it has been determined that decoding can be terminated after $2g$ or $3g$ iterations. To avoid wasting operations on such cases, row operations can first be performed on the coding vector and then repeated on the coded symbol [22]. In both cases, the overhead arises because the symbol is a linear combination of already received symbols.

A simple optimization in cases where the w of the incoming packet is lower than the w of the held symbol with the same pivot candidate, is to simply swap these two to guarantee that w is never increased. Our current implementation does not support this and we leave it to future work to test whether this increases the decoding throughput. However, previous experiments showed that such optimizations can introduce a high cost in terms of bookkeeping [22].

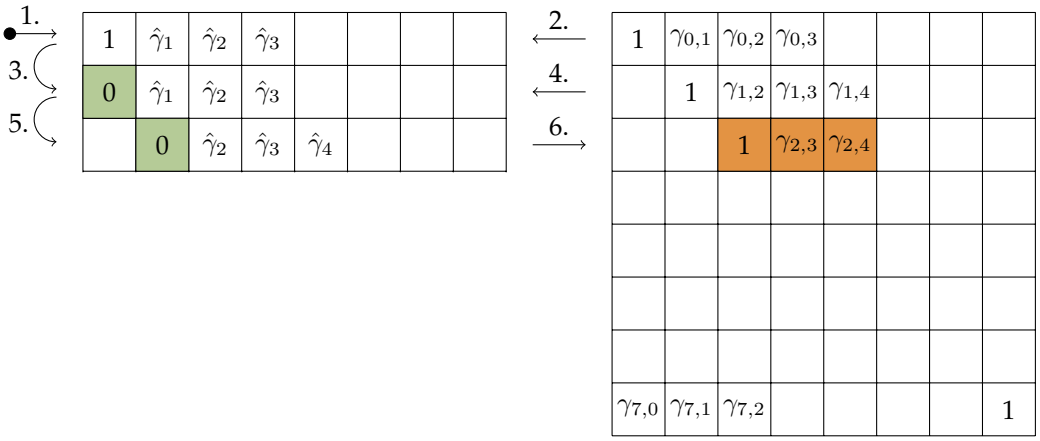

Fig. 3: *On-the-fly* decoding of a received coded packet. The right hand-side matrix is the decoding matrix $\hat{\boldsymbol{G}}$. The left hand-side matrix shows the coding vector of the incoming symbol as it is decoded. The γ's denotes random field elements. The filled circle and arrow indicate the original coding vector of the incoming packet. The straight lines indicate which rows are substituted into the coding vector. The arcs indicate the decoding steps.

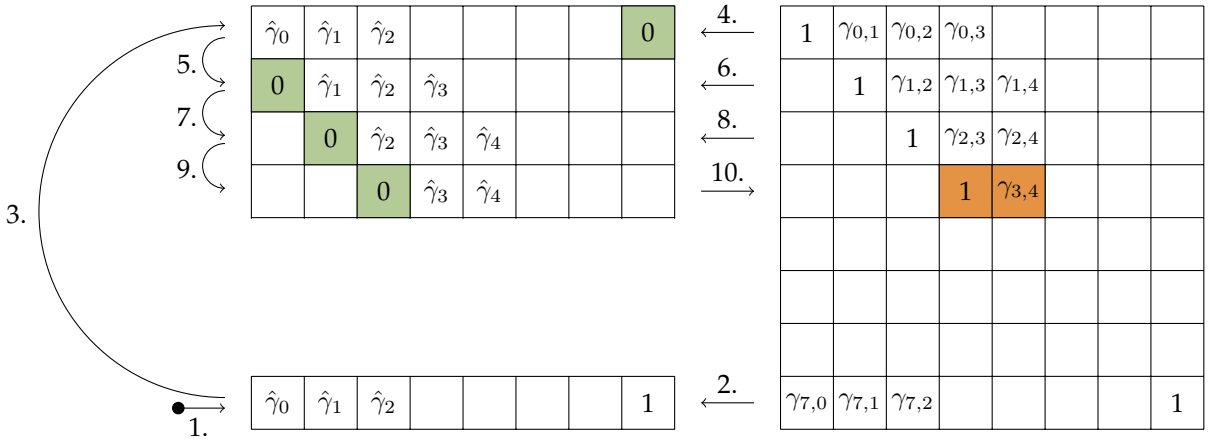

Fig. 4: *On-the-fly* decoding similar to Fig. 3, but the pivot candidate wraps around the end of the decoding matrix.

3.2.2 Final Decoding

When a pivot candidate has been identified for all rows, final decoding is performed by forward substitution and backwards substitution. Initially, the decoding matrix has a form similar to that shown in Fig. 5a. Note that some of the elements $\gamma_{i,j}$ might be zero. It should also be noted that even though a pivot candidate has been identified for all rows, this does not guarantee that the decoding matrix has full rank. Therefore it is important to perform the final decoding in a way that ensures that the decoding matrix is not left in a state where future decoding becomes impossible or problematic.

To bring the matrix onto echelon form, forward substitution is performed on the non-zero elements in the lower left corner of Fig. 5a. When forward substitution is performed on the first column, non-zero elements can be introduced in the lower w rows and further substitution becomes necessary as illustrated on Fig. 5b. After the forward substitution step, the decoding matrix is brought onto echelon form in Fig. 5c.

The final forward step in Algorithm 3 ensures that the decoding matrix is always left in a valid state even if partial final decoding occurs. This happens in cases where it turns out that the decoding matrix does not have full rank, even though a pivot candidate for each row was identified.

In Algorithm 3, a pivot element must be defined for each column (i). If no such pivot element can be found, then it means that none of the received symbols can be used to decode the corresponding row, and we need to receive additional symbols. Therefore we traverse all the rows (j) from the diagonal and down, as we know that for all rows above a pivot index have already been identified. When we find a row for which the current pivot element is non-zero, we swap it with the correct row if it is not already at the correct position. We then forward substitute into the rows below. If we iterate to the last row without identifying a pivot element, then we cannot decode the target row, and we discard the symbol that is incorrectly located on row i. However, we do not want to discard useful symbols, therefore we check if the coding vector for row i is equal to the row

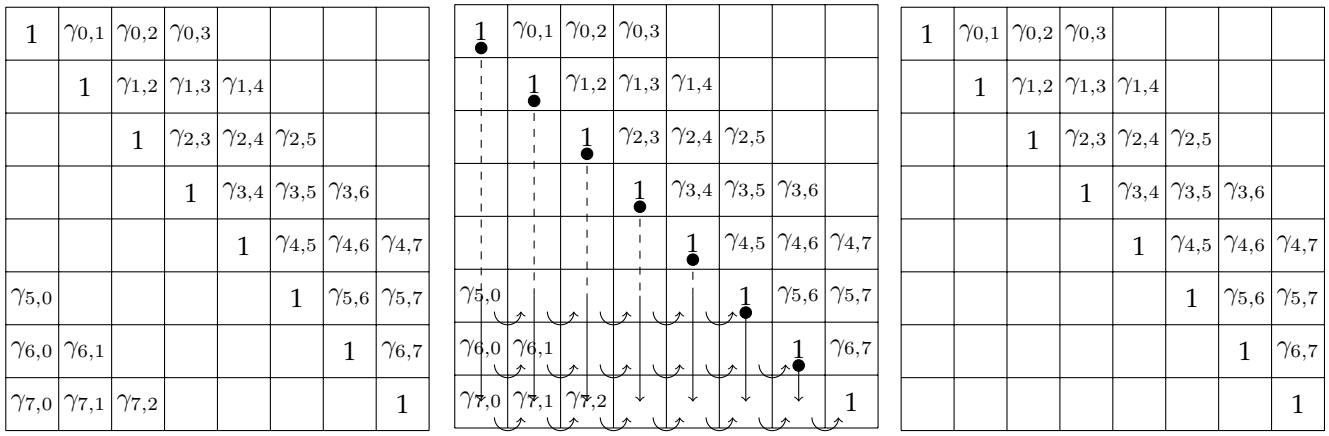

(a) Initial (b) During final forward substitution (c) After final forward substitution

Fig. 5: The decoding matrix $\hat{\boldsymbol{G}}$ at various states of the final decoding. The dotted part of the arrows indicate rows where no substitution is needed. The arching arrows show how the pivot candidate moves towards the diagonal.

Algorithm 3: finalForward

```
for i ∈ [0, ..., g) do
    for j ∈ [i, ..., g) do
        if G_{j,i} ≠ 0 then
            if i = j then
                G_i ← G_i · 1/G_{i,i}
                X_i ← X_i · 1/G_{i,i}                ▷ normalize
            if i ≠ j then
                G_i ↔ G_i · 1/G_{j,i}
                X_i ↔ X_i · 1/G_{j,i}                ▷ normalize and swap
            for k ∈ [max(j + 1, g − w), g) do
                if G_{k,i} ≠ 0 then
                    G_k ← G_k ⊕ G_i
                    X_k ← X_k ⊕ X_i                  ▷ substitute down
            break                                    ▷ found a pivot, skip to the next
        else
            if j = (g − 1) then
                if i ≠ (g − 1) then
                    if G_i ≠ G_{i+1} then
                        G_{i+1} ← G_{i+1} ⊕ G_i
                        X_{i+1} ← X_{i+1} ⊕ X_i
                                                     ▷ add to below symbols
                G_i ← 0
                X_i ← 0                              ▷ discard symbol
```

below, if not we simply add it to the row below. If the two symbols were identical, the result would be the **0** vector and we would discard two rows instead of one. Finally, if we are looking for the pivot element for the final column, there are no rows below our target row, and we simply discard without checking. In this way, the decoding matrix will always be brought as close to echelon form as possible.

If the rank of the matrix is full after the forward substitution, then standard backward substitution is performed to bring the decoding matrix to reduced echelon form and decode the original data.

Algorithm 4: finalBackward

```
for i ∈ (g, ..., 0) do
    for j ∈ (i, ..., max(i − w, 0)] do
        if G_{j,i} ≠ 0 then
            G_j ← G_j ⊕ G_i · G_{j,i}
            X_j ← X_j ⊕ X_i · G_{j,i}
```

Starting from the bottom, all rows are used to remove any remaining non-zeros in the rows above. Note that for each index, it is only necessary to inspect the above w rows as all other rows are guaranteed to be zero due to the form of the decoding matrix. Algorithms 2-4 can be combined to create the decoder in Algorithm 5.

Algorithm 5: decode

```
Input: g, x
if forwardSubstitute(g, x) ≠ −1 then
    if rank(G) = g then
        finalForward()
    if rank(G) = g then
        finalBackward()
return rank(G)
```

When a new packet arrives, it is first forward substituted. If a new pivot element is identified, the coding vector and the coded symbol are inserted into the decoding matrix. When the rank of the decoding matrix is full, final decoding is attempted using forward substitution. This might initially fail, but when it succeeds final

backwards substitution is performed and the original data in the generation is decoded.

3.3 Recoding

When two or more coded or non-coded symbols have been received, they can be combined by recoding. This can be described by Equation (4) and (5) where the collected coding vectors and coded symbols are combined as defined by $\boldsymbol{h}$ of length g', where g' is the number of received symbols. Then $\tilde{\boldsymbol{x}}$ and $\tilde{\boldsymbol{g}}$ together form a recoded packet.

$$\tilde{\boldsymbol{g}} = \hat{\boldsymbol{G}} \cdot \boldsymbol{h} \tag{4}$$

$$\tilde{\boldsymbol{x}} = \hat{\boldsymbol{X}} \cdot \boldsymbol{h} \tag{5}$$

In classical RLNC, coded packets are accumulated and recoding is performed as a separate operation which results in a significant computational load, we denote this type of recoding *active* recoding. As explained in [21], this form of recoding is not suitable when the code is sparse, because the recoded symbol will become denser with high probability [6], [22]. To combat this problem, we introduce a new type of recoding called *passive* recoding.

3.3.1 Active Recoding

Combining all collected packets completely at random, as in standard RLNC, results in recoded packets where the non-zero elements are no longer confined to w elements. If we instead pick packets that have similar pivot elements, then in the worst case the resulting coded packet will only have slightly more non-zero elements w' than that of the original coding vectors. This decreases the freedom in recoding, but allows us to maintain the sparsity in recoded packets. Unfortunately, such an approach significantly increases the complexity of recoding as it introduces a search for an appropriate set of coding vectors. Additionally, it is more deterministic than the standard recoding approach, and thus great care must be taken to avoid generating more linearly dependent symbols.

3.3.2 Passive Recoding

When *on-the-fly* decoding is performed, previously received symbols are subtracted from an incoming symbol to partially decode it. This combining of packets can also be considered as recoding and therefore the operations can be reused in order to reduce the computational load of recoding.

If the operations performed on the received symbols are tracked, a symbol where a sufficient number of operations have been performed can be used as a recoded symbol. One way is to keep a list for each received symbol, to record what symbols are substituted into the symbol. However, this could become unfeasible if g is high. It is simpler to hold an integer for each symbol that is used to count the number of other symbols that have been substituted into the symbol. It is important to remember that during decoding we attempt to decode the symbols, therefore symbols that have been reduced *too much* should not be used as recoded symbols directly. We note that this *passive* approach can also be used for other codes.

3.3.3 Active plus Passive Recoding

To combine the two types of recoding we can monitor the passive recoding. If some neighboring set of packets combined meet our criteria for row operations, we can combine these by *actively* recoding them and thus obtain a recoded symbol. With this hybrid approach, we can recode symbols whenever we need them and still reduce the computation work associated with recoding.

3.3.4 Re-encoding

When a receiver has decoded a generation, it can encode packets the same way as the original source. This is sometimes referred to as recoding, which we believe is misleading, and instead denote this re-encoding to distinguish this from encoding at the original source.

3.3.5 Implementation

In our implementation we have chosen to implement a simple version of the *active plus passive recoding*. The primary reason is to reuse the operations performed during decoding and at the same time allow recoding to be performed when and as much as desired. Another important consideration is to avoid introducing a deterministic behavior when recoding.

Algorithm 6: recode

Input: $\hat{\boldsymbol{G}}, \hat{\boldsymbol{X}}$
if $\operatorname{rank}(\hat{\boldsymbol{G}}) = 0$ **then**
  return -1 ▷ no symbols available
$p \leftarrow (? \mod g)$
while $\hat{\boldsymbol{G}}_{p,p} = 0$ **do**
  $p \leftarrow (? \mod g)$ ▷ find pivot index
$\boldsymbol{h}_p \leftarrow 1$
for $i \in (p, p + w]$ **do**
  if $\hat{\boldsymbol{G}}_{i,i} \neq 0$ **then**
    $\boldsymbol{h}_{(i \mod g)} \leftarrow (? \mod q)$ ▷ draw w elements
$\tilde{\boldsymbol{g}} \leftarrow \hat{\boldsymbol{G}} \cdot \boldsymbol{h}$
$\tilde{\boldsymbol{x}} \leftarrow \hat{\boldsymbol{X}} \cdot \boldsymbol{h}$ ▷ perform recoding
return $\boldsymbol{g}, \boldsymbol{x}$

First, row indices are drawn at random until a row that is non-zero is identified. Then for each of the following w rows that are non-zero, a random coefficient is drawn which defines the recoding vector $\boldsymbol{h}$. The remaining indices in $\boldsymbol{h}$ are zeros. The new coding vector and coded symbol are then computed as $\tilde{\boldsymbol{g}} = \hat{\boldsymbol{G}} \cdot \boldsymbol{h}$ and $\tilde{\boldsymbol{x}} = \hat{\boldsymbol{X}} \cdot \boldsymbol{h}$, respectively.

This approach ensure that the width of the recoded vector $w' \leq 2w$. However, it is worth observing that if a symbol with a higher w is received, the size of w can be reduced during the forward substitution. This would happen more frequently if the width of the received and the existing row is compared as mentioned in Section 3.2.

4 Analysis and Experiments

In this section we present analytical and experimental results on the code overhead, complexity and throughput. To verify the analytical expressions, we have implemented the proposed code in C++ [30]. This also provides us with the possibility to report on encoding and decoding throughput which is a more interesting parameter that defines the computational load at the coding nodes. The current implementation is well tested and we believe that it provides a good trade-off between simplicity and throughput. As the code is available under a research friendly license, we encourage suggestions that can improve the throughput or simplify the implementation.

4.1 Overhead

The code overhead depends on the field size, density, generation size and possibly other factors. From standard RLNC, we have a lower bound for the code overhead as defined in Equation (6), see [21]. The same lower bound holds here, as the lowest overhead is obtained when $w = g - 1$, in which case the perpetual code becomes identical to RLNC.

Equation (6) evaluates the expected overhead based on the probability that the rank increase at the receiver when a new coded symbol is received. This is a function of the generation size, g, the field size, q, and the rank at the receiver, g'. For each of the indices where the decoder has already identified a pivot element, the coefficient in the incoming packet is reduced to zero by the decoder. In the best case, the remaining $g - g'$ elements can be considered as drawn at random from $\mathbb{F}_q$. Hence the probability that these are all zero and the packet is linearly dependent is $1/q^{g-g'}$. The mean overhead is then calculated as the sum of the expected amount of overhead for the decoding of each packet, for all possible ranks of the decoder. Note that the overhead is primarily due to the last packets, and that it becomes negligible for high values of q.

$$\begin{aligned}\alpha &\geq \sum_{g'=0}^{g-1} \left(\left(1 - \frac{1}{q^{g-g'}} \right)^{-1} - 1 \right) \\ &= \sum_{g'=0}^{g-1} \left(\frac{1}{q^{g-g'} - 1} \right)\end{aligned} \tag{6}$$

For a symbol to be independent, either its pivot or one of the w coefficient must hit a new pivot element. The pivot of the symbol is independent and hence the probability is $\frac{1}{g}$, but the w elements depend on the pivot. The probability that one of these w elements hits an uncovered pivot is $\frac{r'}{g}$ where $r' = [1, g-1]$. The expected number of tries to hit an unseen pivot is thus $\sum_{r'=g}^{1} \left(\frac{r'}{g}\right)^{-1} = g \cdot \sum_{r=0}^{g-1} \frac{1}{g-r}$. Thus the probability that one of the w elements hits an unseen pivot can be expressed as $w / \left(g \cdot \sum_{r=0}^{g-1} \frac{1}{g-r} \right)$. Then the probability that a symbol is covered when x coded symbols have been received can be found as the probability that none of the x coded symbols covers the symbol. In the worst case, decoding is possible when all g pivots are covered.

$$F_X(x) \geq \left(1 - \left(1 - \left(\frac{1}{g} + w / \left(g \cdot \sum_{r=0}^{g-1} \frac{1}{g-r} \right) \right) \right)^{x} \right)^{g} \tag{7}$$

The resulting *cdf* can be used to calculate an upper bound for the code overhead by evaluating the corresponding *survival function (sf)*, which defines the probability that there is an uncovered symbol after x transmissions and thus additional transmissions are necessary.

$$\begin{aligned}\beta &\leq \sum_{x=g}^{\infty} S_X(x) = \sum_{x=g}^{\infty} 1 - F_X(x) \\ \alpha &\leq O \leq \alpha + \beta\end{aligned} \tag{8}$$

In Fig. 6, the overhead for different generation sizes is plotted as a function of the code width (shown on the x-axis). The resulting overhead is given on the y-axis is. The dotted lines denote the upper and lower bounds, respectively.

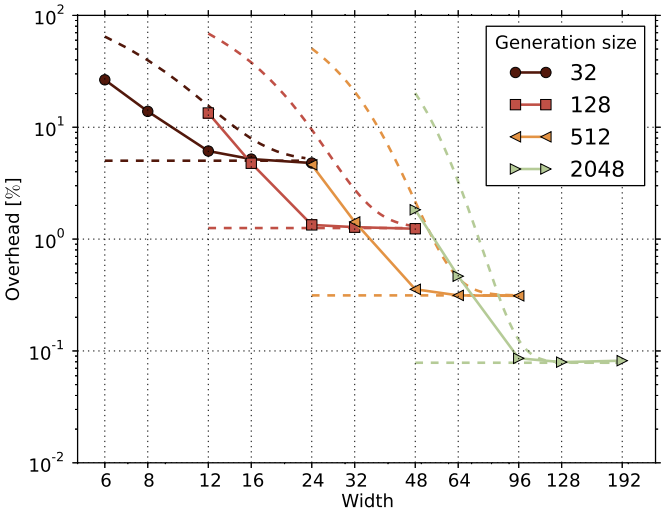

Fig. 6: Code overhead as a function of g and w. The dotted lines denote the lower and upper bounds, respectively.

For each generation size, the overhead decreases as the width increases until the width is sufficiently high and the overhead becomes indistinguishable from the lower bound. If the width is decreased below the sufficient

level, the overhead increases significantly. Therefore, values of w below this point should generally not be used. The bounds are loose for low values of w, but become tighter as w increases. Thus the provided bounds are useful for identifying a value of w that is sufficiently high.

We note that these results do not follow the overhead as a function of the density defined in [21], which is similar to the width for this code. This is not surprising as the code investigated here is significantly less random compared to the sparse RLNC considered in the reference.

4.2 Complexity

We express the computational complexity using a compound metric called row *multiplication-addition*, where a multiplication-addition is multiplying a row with a scalar and adding or subtracting it to or from another row. Here we only consider the binary field, therefore the multiplication scalar is always one and a row multiplication-addition is simply adding or subtracting a row to or from another row.

To *encode* a single packet, the expected number of row operations is given by Equation (9). We start with an empty vector and first add the chosen pivot row to it. For each of the following w indices, the corresponding row is multiplied with a random element from $\mathbb{F}_q$ and added to the new row. The probability that a randomly drawn element from $\mathbb{F}_q$ is non-zero is $1 - \frac{1}{q}$.

$$1 + w \cdot (1 - \frac{1}{q}) \tag{9}$$

During *on-the-fly* decoding, forward substitution is performed on the incoming symbol until a new pivot candidate is identified or the symbol is reduced to the $\mathbf{0}$ vector. Forward substitution continues as long as the next element is non-zero, probability $1 - \frac{1}{q}$, and has not already been identified as a pivot, probability $1 - \frac{r}{g}$, where r is the current rank of the decoding matrix. The expected number of row operations for a generation is found by summing over the reciprocal for all values of r, from which the expected number of operations per symbol is found by dividing with g.

$$\begin{aligned}
\delta_{\text{fly}} \leq & \frac{1}{g} \sum_{r=0}^{g-1} \left(\left(1 - \frac{1}{q}\right) \left(1 - \frac{r}{g}\right) \right)^{-1} \\
= & \frac{q}{q-1} \cdot \frac{1}{g} \sum_{r=0}^{g-1} \frac{g}{g-r} \\
= & \frac{q}{q-1} \sum_{r'=1}^{g} \frac{1}{r'}
\end{aligned} \tag{10}$$

For the upper bound for the *final decoding*, we consider the worst case where most scalars are non-zero, see Fig. 5a. The final forward stage on Fig. 5b can be considered in two steps. First, the bottom w rows are reduced, by substituting the top $g-w$ into them, so only the last w elements are non-zero, hence Equation (11). Then the bottom w rows are brought onto echelon form. Equation (12) accounts for the forward substitution step in the bottom right $w \times w$ submatrix. To include the probability that an element in $\mathbb{F}_q$ is equal to zero, we multiply with $\left(1 - \frac{1}{q}\right)$ and divide by g to find the operations per packet, which is rewritten as $\left(\frac{q-1}{q \cdot g}\right)$.

$$\delta_{\text{forward1}} \leq \left(\frac{q-1}{q \cdot g}\right) \cdot (g - w) \cdot w \tag{11}$$

$$\begin{aligned}
\delta_{\text{forward2}} \leq & \left(\frac{q-1}{q \cdot g}\right) \cdot \sum_{i=1}^{w-1} i \\
= & \left(\frac{q-1}{q \cdot g}\right) \cdot \frac{w \cdot (w-1)}{2}
\end{aligned} \tag{12}$$

To finalize the decoding, a similar procedure is performed, but this time upwards. Each of the $g-w$ bottom rows are substituted into the w rows directly above them. Thus the number of operations is exactly the same as in Equation (11) and Equation (12) and we obtain Equation (13).

$$\begin{aligned}
\delta \leq & \delta_{\text{fly}} + 2\delta_{\text{forward1}} + 2\delta_{\text{forward2}} \\
= & \frac{q}{q-1} \sum_{r'=1}^{g} \frac{1}{r'} + \left(\frac{q-1}{q \cdot g}\right) (w \cdot (2g - w - 1))
\end{aligned} \tag{13}$$

Fig. 7 shows the upper bound and measured number of row multiplication-additions performed to decode one generation, both during the *on-the-fly* and *final* decoding phase. The operations during the two phases are stacked to show the total number of row operations.

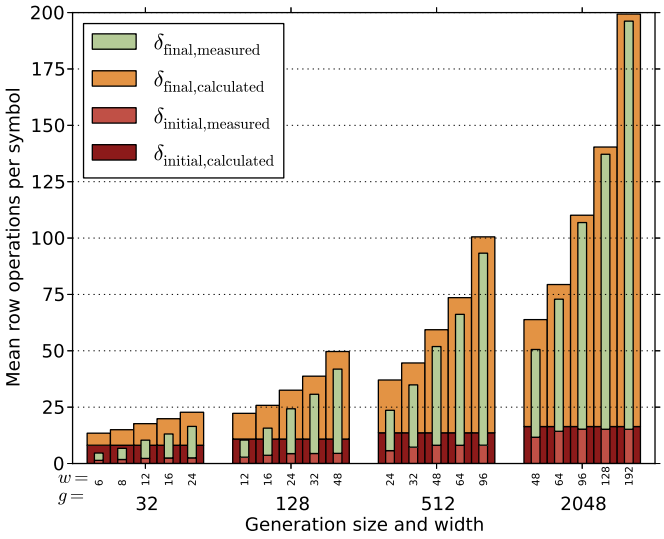

Fig. 7: Mean row operations per decoded symbol.

The analytical expressions for the *on-the-fly* and *final* decoding provide good bounds for the measured results, especially when w is sufficiently high. For low values of w, the bound is less tight, but such settings should not be used when the code overhead is considered.

These values can be compared with traditional RLNC where the expected number of operations to decode a packet is approximately $g/2$ for the binary case [22]. Thus the reduction in complexity compared to RLNC grows as g increases.

4.3 Throughput

A low complexity does not guarantee a low computational load and therefore we investigate the coding throughput. This is due to the complexity introduced by the algorithms that determine how decoding is performed, the quality of their implementation and the bookkeeping they add. The architecture also affects the throughput due to cache misses, memory delay and throughput, which can also be influenced by the memory access pattern.

For each setting, data was encoded and subsequently decoded on the same machine. Each setting was run for a minimum of 30 minutes to reduce the deviation.

TABLE 3: Specifications of the test machine.

Model	Dell Optiplex 790 DT
CPU	Intel Core i7-2600 @ 3.40GHz, 8192 KB L2 cache
Memory	16GB DDR3 1333 MHz Dual channel
Chipset	Intel Q65 Express
OS	64bit Debian Wheezy
Compiler	GNU G++ 4.6

To provide a comparison, we have performed benchmarks of our RLNC implementation using the same values of g as for the perpetual code. The encoding and decoding throughput's are listed in Table 4 along with the corresponding gains over RLNC.

As expected, the throughput for both encoding and decoding decreases as g and w increase. The gain increases for higher values of g which corresponds with the analytical results. The highest gain in encoding and decoding throughput is observed at the highest tested generation size of 2048, and approximately eleven and nine times that of RLNC respectively. It should also be noted that using an excessively high w should be avoided as it decreases the throughput without reducing the code overhead. Additionally, a higher w increases the size of the coding vector representation, see Equation (2), which adds to the overall overhead.

To make a fair comparison with RLNC we must consider both the overhead and the complexity / throughput simultaneously, as the performance of the perpetual code is a trade-off between overhead and speed. As the lower bounds are the same as for RLNC we can never hope to achieve a lower code overhead. However, we can achieve a similar overhead but at lower computational complexity. For this reason the throughput's for the perpetual approach is marked for the lowest value of w where the code overhead is similar to RLNC.

Ideally, all decoding should be performed *on-the-fly* as this decreases the final decoding delay and distributes the processing load evenly. At the same time, decoding should be performed in such a way that *fill-in* does not occur as this reduces the amount of work necessary to decode [6]. In our presented results, the ratio of operations performed during *on-the-fly* phase is low, see Fig. 7. Fortunately, the structure of the code makes it possible to perform something that can best be described as opportunistic backwards substitution. Our tests with this approach show that most of the decoding operations can be performed when symbols are received. However, this algorithm is more difficult to analyze and due to space constraints we have omitted it.

We note that the implementation does not take advantage of multiple cores. This could however be exploited by encoding multiple streams simultaneously or encoding simultaneously from different blocks of the same data set.

TABLE 4: Measured encoding and decoding throughput for RLNC and the proposed perpetual code. Throughputs are reported for different generation sizes, g, and in the case of the perpetual code at different widths, w. The lowest tested value of w where the perpetual code provides a similar code overhead as RLNC is marked.

g	w	Overhead [packets]	Encoding [MB/s]	Gain [%]	Decoding [MB/s]	Gain [%]
32	6	8.24	2883.91	156	2879.99	147
	8	4.15	2512.14	123	2034.45	74
	12	1.70	1915.58	70	1359.66	16
	16	1.65	1506.90	34	1214.37	4
	24	1.62	1080.19	-4	881.03	-25
	RLNC	1.61	1126.16	-	1167.67	-
128	12	17.06	1612.36	441	1209.09	326
	16	6.05	1309.09	339	951.04	235
	24	1.65	951.04	219	620.30	119
	32	1.64	742.68	149	526.43	86
	48	1.63	520.06	74	360.34	27
	RLNC	1.61	298.28	-	283.69	-
512	24	24.33	747.29	921	490.50	655
	32	7.02	612.69	737	392.39	504
	48	1.68	449.39	514	273.00	320
	64	1.65	354.48	384	228.40	251
	96	1.63	249.25	241	167.27	157
	RLNC	1.61	73.17	-	64.99	-
2048	48	36.01	314.79	1656	203.11	1321
	64	9.03	263.36	1369	167.82	1074
	96	1.66	198.93	1010	129.93	809
	128	1.64	160.38	795	99.90	599
	192	1.62	115.00	541	66.81	368
	RLNC	1.61	17.93	-	14.29	-

4.4 Recoding

In general, the performance of the proposed recoding approach will depend on the network topology, therefore general results are difficult to obtain. Instead, we consider the simplest multi-hop topology to provide fundamental insights into the recoding performance of the scheme.

Source A transmits data to R with some erasure probability. Both R and B send a single bit of feedback,

namely when they have achieved full rank. When a symbol arrives at R, it is forwarded to B and to correct the ϵ_{RB} erasures on average $(1-\epsilon_{RB})^{-1}-1$ symbols are recoded at R and transmitted to B. Initially, when too few symbols have been accumulated at R, it is pointless to attempt recoding, which is also the case for traditional RLNC. Therefore, the remaining missing symbols will be re-encoded after R has achieved full rank.

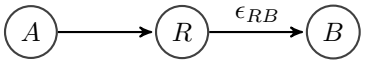

Fig. 8: A simple multihop scenario.

We define the parameter $1 < \mu < g$ which specifies the minimum number of symbols that should be combined to create a recoded symbol. In traditional RLNC, typically all received symbols are combined at random, thus $\mu \approx r \cdot (1 - \frac{1}{q})$. [1] When R combines μ symbols to create a recoded symbol, the probability that one or more symbols previously not seen at B is included in the recoded symbol can be expressed as in Equation (14).

$$P_{\text{unseen}} \leq 1 - (1 - \epsilon_{RB})^{\mu} \tag{14}$$

For $\epsilon = 0.3$ [31] and $P_{\text{unseen}} \approx 0.99$, $\mu = 12$. When a new symbol is received at R, decoding is attempted and with some probability μ or more row operations are performed on the symbol in which case the resulting inserted symbol has been successfully *passively* recoded. This probability, P_{passive}, is found as the probability that a random sequence of μ symbols is non-zero, where the probability of each symbol can be calculated from r and g.

$$\begin{aligned} P_{\text{passive}}(r,\mu) &= \frac{r}{g} \cdot \frac{r-1}{g-1} \cdot \ldots \cdot \frac{r-(\mu-1)}{g-(\mu-1)} \\ &= \prod_{i=0}^{\mu-1} \frac{r-i}{g-i} \end{aligned} \tag{15}$$

Otherwise *active* recoding becomes necessary, and is possible if μ pivot elements have been identified in some range of size Δ. The symbols in this range are combined at random and the resulting symbol will have a width $w' \leq w + \Delta$. The maximal width accepted is denoted $w'_{\max} = w + \Delta_{\max}$. Here we assume that $\Delta_{\max} = 2 \cdot \mu$ in order to permit some freedom during recoding.

Consider a range of size Δ, the number of ranges where at least one pivot element is zero is defined as $\sum_{j=0}^{\mu-1} \left(\binom{\Delta}{j} \binom{g-\Delta}{r-j} \right)$. From this and the total number of combinations $\binom{g}{r}$, the probability that a range contains μ or more pivots can be found. As there are g such ranges we can find the probability that at least one range is suitable, see Equation (16).

$$P_{\text{active}}(r,\mu,\Delta) = 1 - \left(\sum_{j=0}^{\mu-1} \left(\binom{\Delta}{j} \binom{g-\Delta}{r-j} \right) / \binom{g}{r} \right)^{g} \tag{16}$$

The rest of recoding is performed as re-encoding after B has obtained full rank.

$$\begin{aligned} P_{\text{re-encode}}(r,\mu,\Delta_{\max}) = 1 - P_{\text{passive}}(r,\mu) - \\ \sum_{\Delta=\mu}^{\Delta_{\max}} P_{\text{active}}(r,\mu,\Delta) \end{aligned} \tag{17}$$

Finally, we sum over $r = [1, g]$ for P_{passive}, P_{active}, and $P_{\text{re-encode}}$ to obtain the distribution of the recoded output symbols. Passive recoding is preferred over active recoding, and active recoding over a smaller range is desirable. The resulting distribution is illustrated in Fig. 9. The x-axis denotes the maximum Δ and the y-axis denotes the expected ratio of recoded packets. Probabilities for passive recoding, active recoding, and re-encoding are shown.

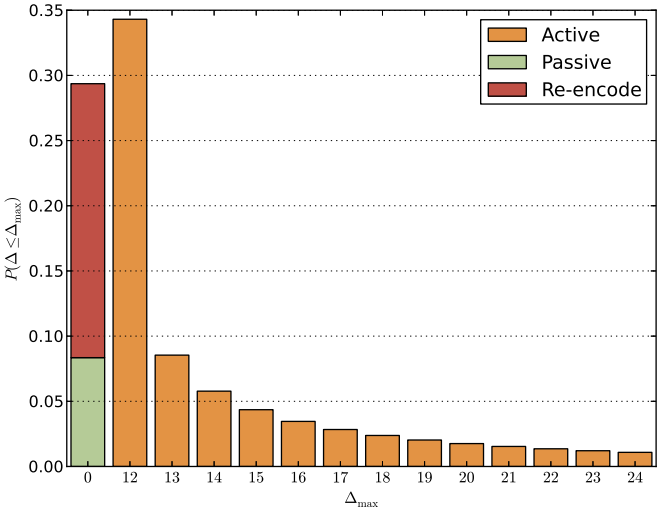

Fig. 9: Distribution of Δ for generated recoded symbols for $g = 512$, $w = 48$, $\epsilon_{RA} = 0.3$, $\mu = 12$, and $\Delta_{\max} = 2 \cdot \mu$.

For 30% of the symbols, whereof 20% is re-encoded, $w' = w$ and the generated symbols are indistinguishable from symbols encoded at a source. For a third of the generated symbols $w' = 48 + 12 = 60$ is slightly larger than $w = 40$. The remaining generated symbols have w' in the range [61,72]. This demonstrates that most times a recoded symbol can be generated and that the expected Δ is low. This is important in order to insure a low delay at R and fast decoding at B.

For lower g, recoding becomes more difficult as w' approaches g. However, more advanced decoding algorithms can be employed to reduce the added Δ, unfortunately space does not permit the inclusion of these.

1. In general, the necessary amount of recoding depends on the network topology, and more specifically on the correlation of incoming links at nodes that receive recoded symbols.

5 Conclusion

In this paper, we presented our initial findings on *perpetual codes* which are suitable for RLNC. We described how encoding, decoding, and recoding can be performed and listed the necessary algorithms. We provided initial analysis of the code performance in terms of overhead and complexity. The analytical results were compared with measurements obtained from our C++ implementation from which we also obtained coding throughput measurements.

The analysis and the tests showed that the proposed approach can obtain a coding overhead similar to RLNC, but at a much lower computational cost. For all tested settings resulting in a code overhead similar to that of RLNC, the proposed approach led to improved encoding and decoding throughput. For the highest tested generation size of 2048, the decoding throughput was almost one order of magnitude higher than that of RLNC. Additionally, the approach provides an easily adjustable parameter that allows for a trade-off between coding complexity and code overhead.

Throughout this work, we have compared the proposed perpetual code with RLNC and SRLNC and not with other rateless FEC codes. The reason is that NC codes allow for recoding where traditional FEC codes do not, thus they are less suitable for cooperative networks. Compared to RLNC and SRLNC, the perpetual code provides the following benefits depending on the chosen values of g and w:

1) Faster encoding, recoding, and decoding.
2) Sparsity is retained when recoding.
3) Small coding vector representation.
4) Simple decoding algorithms.

As the code is sparse, fast encoding is trivial. Fast decoding is possible due to the structure of the code that helps to avoid *fill-in* during decoding. Recoding can be performed fast by using the suggested *passive* recoding approach. We note that this trick can also be employed for other NC codes.

The structure and density of the coded packets can be retained if recoding is performed more carefully than what has been proposed for standard RLNC, we have denoted this *active* recoding. We note that doing so limits the degrees of freedom when recoding, but we believe that our proposed *passive plus active* recoding presents a good trade-off between these two approaches.

As the location of the non-zero elements are well defined, it is trivial to create compact representations of the coding vectors. We believe that this is important as the commonly used assumption of a pseudo random function can be used to compress the coding vector cannot be used if recoding is to be supported [21]. Thus the size of the coding vector must be included in the total overhead.

The presented decoding algorithms are slightly more complicated than for standard RLNC. However, due to the sparsity and structure of the coded symbols, it is possible to eliminate many of the inspections that are necessary when inverting the coding matrix.

We have only considered the *random* encoding mode, meaning that the pivot element is always drawn at random and independently of the previous pivots, see Table 2. This corresponds to the worst-case where the channel is extremely lossy and thus systematic approaches are of no benefit. In cases where the erasure probability is low or moderate, a systematic or sequential mode could be used which would decrease the code overhead and in particular the decoding complexity.

For the future, more rigorous analysis of the code overhead is necessary, especially for the case where low values of w are used. Such analysis would be useful when more advanced variants of the perpetual code are studied. For our implementation, we plan to perform tests using higher field sizes and perform benchmarks on mobile devices. This could help to understand how to choose optimal parameters and to demonstrate the validity of the proposed solution on mobile phones and tablets.

Acknowledgment

This work was partially financed by the CONE project (Grant No. 09-066549/FTP) granted by Danish Ministry of Science, Technology and Innovation.

References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, 2000.

[2] T. Ho, R. Koetter, M. Médard, D. Karger, and M. ros, "The benefits of coding over routing in a randomized setting," in *Proceedings of the IEEE International Symposium on Information Theory, ISIT '03*, June 29 - July 4 2003.

[3] S. Yang and R. W. Yeung, "Large file transmission in network-coded networks with packet loss: a performance perspective," in *Proceedings of the 4th International Symposium on Applied Sciences in Biomedical and Communication Technologies*, ser. ISABEL '11. Barcelona, Spain: ACM, 2011, pp. 117:1–117:5.

[4] J. Heide, M. V. Pedersen, F. H. Fitzek, and T. Larsen, *Network Coding in the Real World*. Academic Press, October 2011, ch. 4, pp. 87–114.

[5] P. Maymounkov. (2006) Perpetual codes: cache-friendly coding. Unpublished draft, retieved 2nd of September 2011. [Online]. Available: http://pdos.csail.mit.edu/~petar/papers/maymounkov-perpetual.ps

[6] S. Ingram. (2006) Minimum degree reordering algorithms: A tutorial. Retrieved March 2010. [Online]. Available: www.cs.ubc.ca/~sfingram/cs517_final.pdf

[7] P. A. Chou, Y. Wu, and K. Jain, "Practical network coding," *Proceedings of the annual Allerton conference on communication control and computing*, vol. 4, pp. 40–49, 2003.

[8] A. Neubauer, J. Freudenberger, and V. Kuhn, *Coding Theory: Algorithms, Architectures and Applications*. Wiley-Interscience, 2007, appendix A provides an in depth review of algebraic structures, including finite fields.

[9] C. Fragouli, J. Boudec, and J. Widmer, "Network coding: an instant primer," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 63–68, 2006.

[10] P. Maymounkov, N. J. A. Harvey, and D. S. Lun, "Methods for Efficient Network Coding," *44th Allerton Annual Conference*, 2006.

[11] Y. Li, E. Soljanin, and P. Spasojevic and, "Collecting coded coupons over overlapping generations," in *Network Coding (NetCod), 2010 IEEE International Symposium on*, june 2010, pp. 1 –6.

[12] O. Trullols-Cruces, J. Barcelo-Ordinas, and M. Fiore, "Exact decoding probability under random linear network coding," *Communications Letters, IEEE*, vol. 15, no. 1, pp. 67 –69, january 2011.

[13] X. Zhao, "Notes on "exact decoding probability under random linear network coding"," *Communications Letters, IEEE*, vol. 16, no. 5, pp. 720 –721, may 2012.

[14] J. Heide, M. V. Pedersen, F. H. Fitzek, and T. Larsen, "Network coding for mobile devices - systematic binary random rateless codes," in *The IEEE International Conference on Communications (ICC)*, Dresden, Germany, 14-18 June 2009.

[15] H. Shojania and B. Li, "Parallelized progressive network coding with hardware acceleration," in *Quality of Service, 2007 Fifteenth IEEE International Workshop on*, June 2007, pp. 47–55.

[16] J. Heide, M. V. Pedersen, F. H. Fitzek, and T. Larsen, "Cautious view on network coding - from theory to practice," *Journal of Communications and Networks (JCN)*, vol. 10, no. 4, pp. 403–411, December 2008.

[17] X. Chu, K. Zhao, and M. Wang, "Massively parallel network coding on gpus," in *Performance, Computing and Communications Conference, 2008. IPCCC 2008. IEEE International*, December 2008, pp. 144–151.

[18] H. Shojania and B. Li, "Random network coding on the iphone: fact or fiction?" in *NOSSDAV '09: Proceedings of the 18th international workshop on Network and operating systems support for digital audio and video*. ACM, June 2009, pp. 37–42.

[19] H. feng (Francis) Lu, "Binary linear network codes," in *IEEE Information Theory Workshop on Information Theory for Wireless Networks*, July 2007.

[20] X. Li, W. H. Mow, and F.-L. Tsang, "Singularity probability analysis for sparse random linear network coding." in *IEEE International Conference on Communications (ICC)*, June 2011.

[21] J. Heide, M. V. Pedersen, F. H. Fitzek, and M. Médard, "On code parameters and coding vector representation for practical rlnc," in *IEEE International Conference on Communications (ICC) - Communication Theory Symposium*, Kyoto, Japan, jun 2011.

[22] J. Heide, M. V. Pedersen, and F. H. Fitzek, "Decoding algorithms for random linear network codes," in *IFIP International Conferences on Networking - Workshop on Network Coding Applications and Protocols (NC-Pro)*, ser. Lecture Notes in Computer Science, vol. 6827, Valencia, Spain, may 2011, pp. 129–137.

[23] E. Erez and M. Feder, "Convolutional network codes," in *Information Theory, 2004. ISIT 2004. Proceedings. International Symposium on*. IEEE, Jun. 2004, pp. 146+.

[24] S. Y. R. Li and R. W. Yeung, "On convolutional network coding," in *Information Theory, 2006 IEEE International Symposium on*. IEEE, Jul. 2006, pp. 1743–1747.

[25] A. Shokrollahi, "Raptor codes," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2551–2567, Jun. 2006.

[26] C. Fragouli and E. Soljanin, "A connection between network coding and convolutional codes," in *IEEE International Conference on Communications (ICC)*, 2004.

[27] S. Jaggi, M. Effros, T. Ho, and M. Mdard, "On linear network coding," in *42st Annu. Allerton Conf. Communication Control and Computing*, 2004.

[28] D. Silva, W. Zeng, and F. Kschischang, "Sparse network coding with overlapping classes," in *Network Coding, Theory, and Applications, 2009. NetCod '09. Workshop on*, june 2009, pp. 74 –79.

[29] Y. Li, E. Soljanin, and P. Spasojevic, "Effects of the generation size and overlap on throughput and complexity in randomized linear network coding," *CoRR*, vol. abs/1011.3498, 2010.

[30] (2012) Perpetual code implementation. Currently not publically available, but will be at time of publication. [Online]. Available: https://github.com/steinwurf/kodo/tree/master/kodo/perpetual

[31] J. Heide, M. V. Pedersen, F. H. Fitzek, T. V. Kozlova, and T. Larsen, "Know your neighbour: Packet loss correlation in ieee 802.11b/g multicast," in *The 4th International Mobile Multimedia Communications Conference (MobiMedia '08)*, Oulu, Finland, July 7-9 2008.

Janus Heide is a postdoctoral researcher at Aalborg University. He received his M.Sc in electrical engineering with specialization Wireless Communication Engineering and his Ph.D. from Aalborg University, Denmark, in 2009 and 2012 respectively. From July 2007 he has been working in the Mobile Devices research group at Aalborg University. His main research interests are protocol analysis and design, network coding, wireless and cooperative communication, and data distribution in meshed networks.

Morten V. Pedersen is a postdoctoral researcher at Aalborg University. He received his M.Sc. in wireless communication in 2009, and his Ph.D. in 2012 both from Aalborg University, Denmark. From 2006 he has been working in the Mobile Devices research group at Aalborg University. In 2010 he was appointed Forum Nokia Champion. His main research interests are mobile programming, cooperative communication, network coding and network performance evaluation.

Frank H.P. Fitzek is a Professor in the department of Electronic Systems, Aalborg University, Denmark heading the Mobile Device group. In 2005 he won the YRP award and received the Young Elite Researcher Award of Denmark. He was selected to receive the NOKIA Champion Award five times in a row from 2007 to 2011. In 2011 he received the SAPERE AUDE research grant from the Danish government and in 2012 he received the Vodafone Innovation price. His current research interests are in the areas of wireless and mobile communication networks, mobile phone programming, cross layer as well as energy efficient protocol design and cooperative networking.

Muriel Médard is a Professor in the Electrical Engineering and Computer Science at MIT. Professor Mdard received B.S. degrees in EECS and in Mathematics in 1989, a B.S. degree in Humanities in 1990, a M.S. degree in EE 1991, and a Sc D. degree in EE in 1995, all from the Massachusetts Institute of Technology (MIT), Cambridge. She has served as editor for numerous journals and on the board of Governors of the IEEE Information Theory Society as well as serving as the first Vice President in 2011 and the President in 2012. Her research interests are in the areas of network coding and reliable communications, particularly for optical and wireless networks.